

Claims

- [c1] 1. A system for allocating resources amongst a plurality of applications , the system comprising:
a plurality of computers connected to one another through a network;
a policy engine for specifying a policy for allocation of resources of the plurality of computers amongst a plurality of applications having access to the resources;
a monitoring module at each computer for detecting demands for the resources and exchanging information regarding demands for the resources at the plurality of computers; and
an enforcement module at each computer for allocating the resources amongst the plurality of applications based on the policy and information regarding demands for the resources.
- [c2] 2. The system of claim 1, wherein the resources include communication resources.
- [c3] 3. The system of claim 2, wherein the communication resources include network bandwidth.
- [c4] 4. The system of claim 1, wherein the resources include

processing resources.

- [c5] 5. The system of claim 1, wherein the resources include a selected one of memory, disk space, system I/O (input/output), printers, tape drivers, and software licenses.
- [c6] 6. The system of claim 1, wherein the policy engine provides receives user input for defining an application subject to the policy.
- [c7] 7. The system of claim 6, wherein the monitoring module identifies an application running at a given computer based, at least in part, upon the user input for defining the application.
- [c8] 8. The system of claim 7, wherein the monitoring module detects a request for resources by the application at the given computer.
- [c9] 9. The system of claim 6, wherein the user input includes defining components of an application.
- [c10] 10. The system of claim 9, wherein the components include a selected one of processes, network traffic, and J2EE components.
- [c11] 11. The system of claim 1, wherein the policy engine receives user input of a policy specifying actions to be

taken for allocation of the resources in response to particular conditions.

- [c12] 12. The system of claim 11, wherein the policy includes a command to be run in response to a particular condition.
- [c13] 13. The system of claim 11, wherein the policy includes an attribute indicating when a particular condition of the policy is to be evaluated.
- [c14] 14. The system of claim 13, wherein the policy includes an attribute indicating when action is to be taken based upon a particular condition of the policy being satisfied.
- [c15] 15. The system of claim 11, wherein the policy specifies priorities of the plurality of applications to the resources.
- [c16] 16. The system of claim 15, wherein the enforcement module allocates resources amongst the plurality of applications based, at least in part, upon the specified priorities.
- [c17] 17. The system of claim 1, wherein the policy engine includes a user interface for a user to specify the policy.
- [c18] 18. The system of claim 1, wherein the policy engine supports an expression language for policy definition.
- [c19] 19. The system of claim 1, wherein the policy engine is a

distributed system operating at each of the plurality of computers.

[c20] 20. The system of claim 1, wherein the monitoring module determines resources available at each computer.

[c21] 21. The system of claim 1, wherein the monitoring module determines resource utilization at each computer.

[c22] 22. The system of claim 21, wherein the monitoring module at each computer exchanges resource utilization information amongst the plurality of computers.

[c23] 23. The system of claim 1, wherein the enforcement module allocates network bandwidth amongst said plurality of applications based upon the policy and information regarding demands for the resources.

[c24] 24. The system of claim 1, wherein the enforcement module allocates processor resources amongst said plurality of applications based upon the policy and information regarding demands for the resources.

[c25] 25. The system of claim 1, wherein the enforcement module includes an interface for communication with an external module for specifying allocation of resources by said external module.

[c26] 26. The system of claim 25, wherein said external mod-

ule includes a load balancer for load balancing instances of an application.

[c27] 27. The system of claim 25, wherein said external module comprises a selected one of a router and a provisioning device.

[c28] 28. The system of claim 1, wherein the enforcement module starts an instance of an application on a given computer based upon the policy and information regarding demands for the resources.

[c29] 29. An improved method for allocating resources of a plurality of computers to a plurality of applications, the method comprising:
receiving user input specifying a dynamically configurable policy for allocating resources of a plurality of computers amongst a plurality of applications having access to the resources;
at each of the plurality of computers, detecting demands for the resources from the plurality of applications and availability of the resources;
exchanging information regarding demand for the resources and availability of the resources amongst the plurality of computers; and
allocating the resources to each of the plurality of applications based on the dynamically configurable policy and

the information regarding demand for the resources and availability of the resources.

- [c30] 30. The method of claim 29, wherein the resources include communication resources.
- [c31] 31. The method of claim 30, wherein the communication resources include network bandwidth.
- [c32] 32. The method of claim 29, wherein the resources include processing resources.
- [c33] 33. The method of claim 29, wherein the resources include a selected one of memory, disk space, system I/O (input/output), printers, tape drivers, load balancers, and software licenses.
- [c34] 34. The method of claim 29, wherein said receiving step includes receiving user input for defining an application.
- [c35] 35. The method of claim 34, further comprising:
providing a set of default rules for assisting a user in defining an application.
- [c36] 36. The method of claim 29, wherein said detecting step includes detecting applications running on each of said plurality of computers.
- [c37] 37. The method of claim 29, wherein said detecting step

associating particular processes with a particular application.

[c38] 38. The method of claim 29, wherein said detecting step associating particular network traffic with a particular application.

[c39] 39. The method of claim 29, wherein said detecting step includes detecting components of an application.

[c40] 40. The method of claim 39, wherein said step of detecting components of an application includes detecting a J2EE component.

[c41] 41. The method of claim 29, wherein said receiving step includes receiving user input specifying actions to be taken for allocation of the resources in response to particular conditions.

[c42] 42. The method of claim 41, wherein the user specifies a script to be run based upon a particular condition.

[c43] 43. The method of claim 41, wherein the user specifies when a particular condition of the dynamically configurable policy is to be evaluated.

[c44] 44. The method of claim 43, wherein the user specifies when action is to be taken based upon the particular condition being satisfied.

[c45] 45. The method of claim 29, wherein said receiving step includes receiving user input specifying priorities of the plurality of applications to the resources.

[c46] 46. The method of claim 45, wherein the allocating step includes allocating resources amongst the plurality of applications based, at least in part, upon the specified priorities.

[c47] 47. The method of claim 29, wherein the receiving step includes providing an expression language for policy definition.

[c48] 48. The method of claim 29, wherein said detecting step includes determining resource utilization at the given computer.

[c49] 49. The method of claim 48, wherein said determining resource utilization step includes determining average resource utilization over a given time period.

[c50] 50. The method of claim 29, wherein said allocating step includes allocating a specified amount of resources to a particular application when the particular application is initially detected at a given computer.

[c51] 51. The method of claim 29, wherein said allocating step includes allocating resources based upon particular

events.

- [c52] 52. The method of claim 29, wherein said allocating step includes allocating network bandwidth to each of the applications.
- [c53] 53. The method of claim 29, wherein said allocating step includes allocating processor resources amongst the plurality of applications.
- [c54] 54. The method of claim 29, wherein said allocating step includes communicating with an external module for allocating resources provided by an external module.
- [c55] 55. The method of claim 54, wherein the external module comprises a load balancer for load balancing instances of an application.
- [c56] 56. The method of claim 54, wherein the external module comprises a selected one of a router and a provisioning device.
- [c57] 57. The method of claim 29, wherein said allocating step includes starting an instance of an application on a given computer.
- [c58] 58. A computer-readable medium having processor-executable instructions for performing the method of claim 29.

[c59] 59. A downloadable set of processor-executable instructions for performing the method of claim 29.

[c60] 60. A method for allocating resources to a plurality of applications, the method comprising:
receiving user input specifying priorities of the plurality of applications to resources of a plurality of servers, the specified priorities including designated servers assigned to at least some of the plurality of applications;
selecting a given application based upon the specified priorities of the plurality of applications;
determining available servers on which the given application is runnable and which are not assigned to a higher priority application;
allocating to the given application any available servers which are designated servers assigned to the given application;
allocating any additional available servers to the given application until the given application's demands for resources are satisfied; and
repeating above steps for each of the plurality of applications based on the specified priorities.

[c61] 61. The method of claim 60, wherein the receiving step includes receiving user input of a value for a given application representing relative priority of the given applica-

tion compared to other applications.

- [c62] 62. The method of claim 60, wherein the receiving step includes receiving a ranking of the plurality of applications from highest priority to lowest priority.
- [c63] 63. The method of claim 62, wherein the step of selecting a given application includes commencing with selection of an application having the highest priority.
- [c64] 64. The method of claim 60, further comprising:
powering on a server allocated to an application if the server is in a powered off state.
- [c65] 65. The method of claim 60, further comprising:
determining whether an application is inactive on a server allocated to the application; and
initiating a resume script for running the application on the server application is determined to be inactive.
- [c66] 66. The method of claim 60, further comprising:
adding a server newly allocated to an application to a set of servers across which the application is load balanced.
- [c67] 67. The method of claim 60, further comprising:
removing a server no longer allocated to an application from a set of servers across which the application is load balanced.

- [c68] 68. The method of claim 60, further comprising:
determining whether a server no longer allocated to an application is in a suspend set of servers designated for the application; and
running a suspend script if the server is determined to be in the suspend set of servers.
- [c69] 69. The method of claim 68, further comprising:
if a suspend script is executed on the server, determining whether the server should be powered off based on consulting a power management rule; and
powering off the server if it determined that the server should be powered off.
- [c70] 70. The method of claim 60, wherein said allocating step includes starting an instance of an application on a given computer.
- [c71] 71. A computer-readable medium having processor-executable instructions for performing the method of claim 60.
- [c72] 72. A downloadable set of processor-executable instructions for performing the method of claim 60.